JOURNAL OF THE EXPERIMENTAL ANALYSIS OF BEHAVIOR 2002, **77,** 3–27 NUMBER 1 (JANUARY)

LINEAR MODELING OF STEADY-STATE BEHAVIORAL DYNAMICS

WILLIAM L. PALYA, DONALD WALTER, ROBERT KESSEL, AND ROBERT LUCKE

JACKSONVILLE STATE UNIVERSITY AND NAVAL RESEARCH LABORATORY

The observed steady-state behavioral dynamics supported by unsignaled periods of reinforcement within repeating 2,000-s trials were modeled with a linear transfer function. These experiments employed improved schedule forms and analytical methods to improve the precision of the measured transfer function, compared to previous work. The refinements include both the use of multiple reinforcement periods that improve spectral coverage and averaging of independently determined transfer functions. A linear analysis was then used to predict behavior observed for three different test schedules. The fidelity of these predictions was determined.

Key words: behavioral dynamics, linear systems analysis, transfer function, quantitative analysis, model, key peck, pigeons

Although the extremely rich repertoires of behavior supported by time-varying schedules of reinforcement have been noted with interest for decades, quantitative models of behavioral dynamics have been lacking. Palya, Walter, Kessel, and Lucke (1996), however, successfully extracted the salient properties for an individual organism from the behavioral dynamics supported by one schedule in a manner that allowed prediction of the behavioral dynamics supported by a second schedule. To a large extent, our progress was made by combining a carefully selected experimental environment with an appropriate data analysis method. The schedules used in our prior report were simple: periodic pulses, or bursts, of reinforcement availability for responses following variable periods with a constant average (a variable-interval or VI schedule), separated by periods during which responses had no effect (extinction). Further, we focused on steady-state responding so the behavioral dynamics themselves were restricted to a particularly simple form. The prediction method used to generalize from one schedule to another was built from an assumed linear dependence of the steady-state behavior upon the supporting reinforcement schedule. Palya et al.'s results not only demonstrated that some forms of behavioral dynamics can be understood but also suggested approaches that could refine this understanding. Our initial improvements are modifications to the experimental procedures, rather than alterations of the underlying linear analysis. How well, how simply, and how broadly we can stretch the assumption of linearity are the core questions that we begin to address in this report. In the longer term, by answering these questions, we may further understand behavioral dynamics.

This report begins with a review of the important stages required when generalizing a specific measurement of behavioral dynamics to make broader predictions. As part of this review, the differences between a linear analysis of steady-state behavioral dynamics and a more traditional analysis of static behavior, such as a matching law analysis, are noted. Next, we cover the refinements suggested by Palya et al. (1996) and introduced in this study. A section describing the experimental procedure is presented next. We then discuss the results obtained from both qualitative

We gratefully acknowledge the contributions of Gregory Galbicka for many helpful discussions during the data reduction and comments on prior drafts of this report and Elizabeth Palya for contributions in all phases of this research. Portions of this paper were presented at the meetings of the Society for the Quantitative Analyses of Behavior, May 1997 and May 1998, and the Southeastern Association for Behavior Analysis, October 1998.

Address comments and other correspondence to Bill Palya or Don Walter, Department of Psychology, Jacksonville State University, Jacksonville, Alabama 36265-1602 (palya@sebac.jsu.edu or walter@sebac.jsu.edu) or Bob Kessel, Electro-Optics Technology Section, Code 8123, Naval Research Laboratory, Washington, D.C. 20375-5354 (kessel@ncst.nrl.navy.mil). The data-reduction software used for the linear analysis based prediction is available from the JSU program archive. The raw data logs are also available upon request from the JSU data archive.

and quantitative perspectives. We close by summarizing the conclusions we have drawn from the results, as well as the questions opened by the results. There is also an appendix that covers the mathematical basis for the data reduction, linear analysis predictions, and testing the quality of the predictions.

GENERALIZING OBSERVED BEHAVIORAL DYNAMICS WITH A LINEAR ANALYSIS

The experiment and linear analysis of this report address specific technical questions: How faithfully can an animal's steady-state periodic boundary behavioral dynamics be described by a linear filter? Are these dynamics invariant properties of the individual pigeon? The selection of steady-state periodic boundary conditions provides the simplest behavioral dynamics problem to address. In steady-state behavior, one can avoid the more complex transient case of adaptive behavioral dynamics in which an animal is picking up and dropping behaviors. We are asserting that in the steady state, when acquisition is complete, the behavioral dynamics should look fairly linear. We further simplify by restricting the problem to that of an averaged steady state. Once the anchor point of this simpler behavioral dynamics has been established, one can then begin work on more complex behavioral models. As a context note to place our usage of the term, we are using *dynamics* to mean time-varying behavior.

Any analysis of behavioral dynamics must describe moment-by-moment changes in behavior over extended periods of time. In fact, the study of how a system's state varies moment by moment over time is, for all intents and purposes, the definition of dynamics. Beyond just describing any specific example of behavioral dynamics, any model, given one set of observed behavioral dynamics, should provide a method that predicts the behavior that will be observed as the organism's environment changes over some reasonable range. The allowed range of changes measures a model's performance; the greater the range, the better the model.

A linear analysis provides one framework that can relate the moment-by-moment stream of behavior from an organism (dependent measure) to the moment-by-moment pattern of reinforcers (independent measure) (McDowell, Bass, & Kessel, 1993). As already noted, Palya et al. (1996) showed that such a linear analysis can, over some range, also satisfy the generalization requirement. McDowell et al.'s form of linear analysis is an inherently dynamic description grouped under the general heading of filter analysis. This ''animal as filter'' concept emphasizes the measured reinforcer and response times rather than hypothetical internal processes, such as the notion that the occurrence of a reinforcer ''strengthens'' a response tendency or strengthens a ''connection'' between a stimulus and a response. It also differs from positions that argue that a molar response rate is a function of a molar reinforcement rate. However, McDowell, Bass, and Kessel (1983) showed that linear analysis has the correct limit in the static regime in which a molar analysis is widely accepted.

Because one of the more common uses of linear analysis is in electrical engineering, an electronic analogy can help to express the character of this approach. For a simple time-dependent schedule, an analogy can be drawn between the pattern of the reinforcers received by the animal over time and the input wave form to an audio amplifier. The output wave form of the amplifier would then be analogous to the behavioral response of the organism. In both cases, the output wave form is seen as being determined by the input wave form. The analogy extends naturally beyond the identification of the input and output to a different means of data description: Both the reinforcement schedule and supported behavior are describable as a summation over frequency components (i.e., sine waves) using the standard methods of Fourier analysis (Bracewell, 1986; Ramirez, 1985). Provided linearity is appropriate, these summations over frequency allow efficient study and prediction of the underlying behavioral dynamics. In the electrical case, the dependence of output on input is purely linear; in the behavioral case, the dependence is not expected to be purely linear, but may have a substantial linear component.

Turning from the electrical analogy, we need to refine the definitions of the elements and relations used for a linear analysis of be-

havioral dynamics. The first element, the pattern of reinforcers, is treated by linear behavior analysis as a complicated signal containing many frequencies: Some reinforcers are close together, some are far apart. In addition, each frequency component has its own respective amplitude, in that some frequencies occur often and some occur rarely. Behavior, the second element, is viewed similarly: Some interresponse times are long, some are short, some are frequent, some are rare. This decomposition of the complex patterns of reinforcers and behavior makes possible more than just a convenient method to catalog the properties of the data. To the extent that the organism satisfies linearity, the change that occurs at each frequency when going from input (reinforcers) to output (behavior) is multiplicative, that is,

$$\text{output}(f) = g(f) \times \text{input}(f). \quad (1)$$

Further, the factor $g(f)$, called the *transfer function,* will be the same for that frequency component of future reinforcements. A perfectly linear organism would be characterized uniquely by a single linear transfer function that is independent of differences in procedures, differences in learning histories, and possibly even differences in concurrently presented contingencies. Equation 1, combined with this independence of $g(f)$, allows generalizing a specific observation to predict the behavioral dynamics expected for new schedules.

The benefits of conceptualizing an organism as a filter are that (a) it makes possible a prediction of the dynamics of behavior rather than some molar or static prediction of a behavioral average. That is, this approach predicts the moment-to-moment changes in the behavior of an individual across some extended period of time. (b) It is based exclusively on the relation of inputs and outputs. A linear analysis is silent on any reductionistic machinery that may produce the filter characteristic of the organism, including the strengthening of molecular or molar behavior by the occurrence of a reinforcer. And (c) the predictions can be made with no free parameters. The importance of this third aspect, with respect to serious theorizing, cannot be overstated. Free parameters often serve a useful purpose in the initial stages of description and analysis, but until they are removed, any model falls short of completeness. Free parameters buy prediction and thereby credibility on credit, but a model cannot be constructed solely of debt. The task of an empirical science is to characterize the factors that control behavior, rather than to accrete free parameters.

How does the purely linear analysis implemented in this report relate to the broader questions of behavioral dynamics? Although a linear analysis is likely to be incomplete or insufficient for a variety of reasons, it is also likely to be a necessary first step in a more general analysis of behavioral dynamics. For example, although the adaptive acquisition of behavior is outside the scope of a linear model, the steady-state behavior may well be an approximately linear limit of adaptive acquisition. Further, to attack questions like adaptive acquisition or isolated schedule transitions with a nonlinear approach, such as the quadratic and higher order terms of the Kubo-Bass series, the range and performance of the linear limit are a necessary starting point (McDowell, Bass, & Kessel, 1992). Starting with a linear analysis also provides access to a class of fully developed analytic and numerical techniques with well-understood properties. Consequently, our experiment focuses on steady-state behavioral dynamics.

Palya et al. (1996) were the first to demonstrate the use of these standard numerical methods in a behavioral dynamics context. The report showed that (a) a linear approximation is sufficient to make reasonable predictions, without free parameters, for a new schedule, and (b) the general qualitative form of pigeons' transfer functions is that of a low-pass filter. Palya et al. also suggested a number of refinements necessary to measure a steady-state transfer function and define more demanding tests of a linear approximation for behavioral dynamics.

Palya et al.'s (1996) determination of a pigeon's transfer function depends on two premises. First, the experimental procedure used to measure the transfer function employs an aspect of Fourier's theorem. In principle, because a linear transfer function is the specification of how a system changes each possible input frequency, every reinforcement frequency must be tested. Doing that individually would be prohibitively time consuming. Fourier's theorem shows that a zero-duration

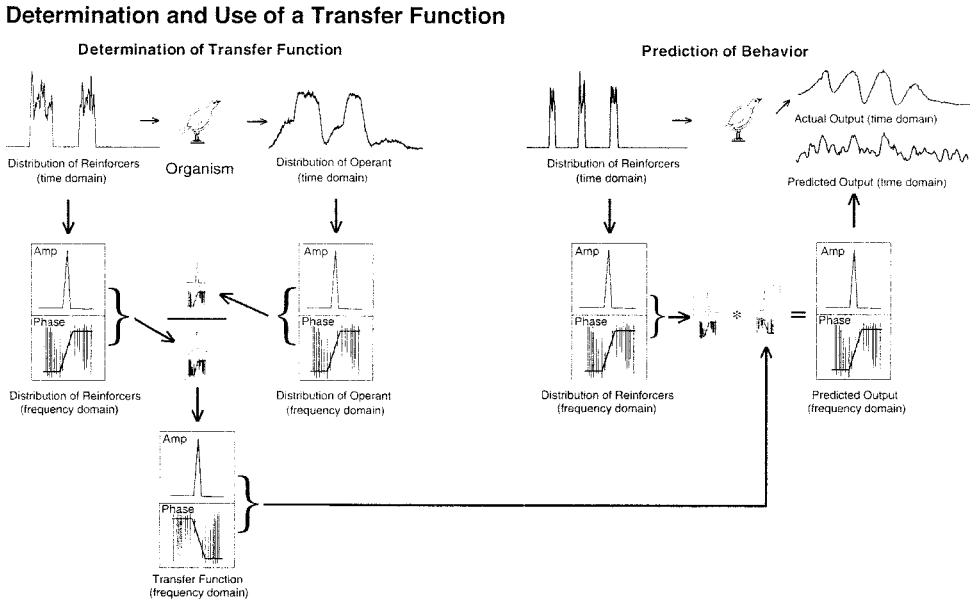

Fig. 1. A schematic representation of the algorithm used, first, to extract a transfer function and then, second, to predict behavior to a new schedule. The upper left images illustrate the behavioral output to a pattern of reinforcers. The middle left depicts the conversion of those data to frequencies and the division of the output by the input. The lower left image represents the transfer function. The left side of the right half of the figure illustrates a schedule and its frequency-domain representation. The lower portion of the right side of the figure illustrates the reinforcement-rate data being multiplied by the transfer function to produce the frequency representation of the predicted behavior. This is transformed to the time domain and compared to the actual behavior, as illustrated in the upper right portion.

pulse can be replicated by an equally weighted sum of all possible frequencies (Bracewell, 1986). In other words, a sum of all possible input frequencies of reinforcement would be obtained by implementing a procedure containing a zero-duration pulse of reinforcers. Lengthening the pulse to a few minutes with a relatively rich reinforcement schedule makes practical implementation possible without an unacceptable loss of frequencies. Palya et al.'s procedure used repetitive exposures to a 1,000-s trial containing an unsignaled step transition from a VI 20-s schedule to extinction after 200 s (a mixed schedule). Second, the numerical computation of a transfer function turns on the assumption of linearity. For any linear system, the transfer function can be calculated by dividing each output frequency component by the input frequency component. Further, the division at each frequency is independent of the division at all other frequencies. The possibility of divide-by-zero errors, however, does have an important consequence: Some reinforcement schedules are more useful than others for estimating a transfer function. We will consider this point further below. The domain change from time-based measures to frequency-based measures is analogous to converting an involved arithmetic problem to logarithms in order to simplify its computation.

Figure 1 illustrates in schematic form the two phases of Palya et al.'s (1996) analysis and the central role played by the transfer function. On the left, the pigeon's dynamics are captured by a transfer function; on the right, the same transfer function is then used to predict behavior for a new contingency. Although each of the transformations shown in Figure 1 is actually done numerically, at an intuitive level, the figure is a complete depiction of the features present in our linear analysis of behavioral dynamics. (The analytic basis of the numerical computations is given by Palya et al. and is briefly repeated in the Appendix. Palya et al. also provided examples of all intermediate results as well as a discussion of units.) The upper left portion illustrates a pigeon emitting some pattern of behavior in

response to some dynamic pattern of reinforcers over some time interval. Below that, in the left central portion, are the frequency domain equivalents of these two time histories. Effectively, these frequency domain equivalents are the spectra of the reinforcement and operant behavior. The conversion produces both amplitude and phase information. Dividing the behavioral spectrum (output) by the reinforcement spectrum (input) yields the transfer function. Note that a consequence of transforming to the frequency domain is that reinforcement, response, and transfer function are all complex-valued functions. Although Figure 1 uses the common and convenient representation of the real and imaginary parts of these three functions as an amplitude and phase, the actual computations done in the frequency domain employ complex arithmetic.

The framework for linear analysis developed by McDowell et al. (1993) relates individual responses to the prior reinforcers. The present experiment's data are not, in a strict sense, used in such a completely local manner within the linear analysis shown in Figure 1. As noted by Palya et al. (1996), some degree of approximation is introduced by the repeated trials local average. A completely local analysis should include explicit bookkeeping that preserves trial order, so each response would be attributed only to prior reinforcement history. A repeated trials local average, in contrast, combines the data from multiple trials and excludes the trial ordering from consideration. The substitution is one of going from an individual trial's reinforcers generating a set of responses to a distribution of reinforcement sequences generating a distribution of response sequences. As also noted by Palya et al., the repeated trials local average can then simplify the behavioral dynamics to the dependence of the response sequences' mean upon the reinforcement sequences' mean. The transfer function that relates the distribution means is not necessarily the same as the mean of the transfer functions determined from each trial individually. Presumably, though, for steady-state observations of behavioral dynamics, it is a reasonable approximation. This approximation will limit the precision with which a transfer function estimate can be used to predict behavior.

The transfer function determined by the processing on the left side of Figure 1 is a characterization of how each possible input frequency is either increased, held constant, or suppressed by the pigeon. At the outset, the transfer function is a set of unknown parameters. Once estimated from data, the transfer function is fixed. This estimation process allows a linear analysis to account for the pigeon-to-pigeon variability. During the second phase of a linear analysis, one predicts the absolute response rate without free parameters. For contrast, if some other type of analysis yielded a functional dependence upon time but left the absolute rate undefined, then the normalization constant required to scale the prediction to the data would be a free parameter.

The second phase of our analysis tests the accuracy of a prediction for behavior controlled by a new pattern of reinforcers. Because the trials are governed by a new reinforcement schedule, the pigeon, after some interval of transient acquisition behavior, emits some new steady-state distribution of behavior. The starting point for the second phase of the analysis is the new distribution of reinforcers illustrated in the upper right portion of Figure 1. The experimental data are used differently in the second phase of the analysis; only the reinforcer distribution is converted to the frequency domain. It is then multiplied by the transfer function that was determined by the first phase of the analysis. The product is the frequency domain representation of the predicted behavior. The prediction is then converted back to the time domain. The prediction has exactly the same form as the measured behavior: the local response rates for all time bins of the test schedule. The time domain form of the prediction, along with the actual behavior supported by the new reinforcer distribution, can be seen in the upper right portion of the figure. The sum of the squared discrepancies between the measured distribution of behavior and the prediction at a uniformly spaced set of times within the trial interval can be assessed with χ^2_ν. Palya et al.'s (1996) comparison of the predictions to chance predictions via χ^2_ν indicated that even the initial implementation of the analysis resulted in significant ($p < .05$) predictions.

In this report we refine the experimental and data reduction procedures by (a) com-

paring the relative efficacy of several data-sampling methods for determining the transfer function; (b) correcting for the incomplete sampling of the input frequency spectrum when determining the transfer function by using a two-pulse procedure; and (c) avoiding the truncation of the output wave form by shifting the initial pulse to later in the interval. The report also sets semiquantitative limits on the precision possible with a linear analysis of behavioral dynamics.

IMPROVED METHODS OF TRANSFER FUNCTION ESTIMATION

The overall performance of a linear analysis of behavioral dynamics depends sensitively on the precision of the transfer function. As noted by Palya et al. (1996), this precision is determined, in part, by the experimental design. Even if the pigeon's behavioral dynamics were completely linear, some reinforcement schedules would yield a better estimate of the transfer function than others. Given that only a preliminary understanding of behavioral dynamics is available, the present experiments employed three refinements of Palya et al.'s procedure that address known difficulties. As more becomes known about the linear component of behavioral dynamics, its characteristics can, in turn, provide the necessary starting point for existing techniques to improve the estimation of a transfer function (Ayers & Dainty, 1988; Fienup, 1993). The iterative techniques would combine an observed property, for example, that the response rates are always greater than or equal to zero throughout the trial, with the initial measured transfer function as the basis of a search algorithm for a refined transfer function that had both the desired property and the best agreement with experimental data. Alternatively, with constraint-based methods, one can introduce a constraint on the allowed transfer function, for example, that it take only forms that can be generated from a time-domain wavelet. Both iterative or constraint-based methods provide a natural framework to include assumed behavioral properties and determine if their inclusion yields an improved linear analysis. These classes of bootstrap processes, however, belong to our field's future. For the present, it is preferable to confine modifications to only the experimental phase durations or the reinforcement schedules, because neither change is likely to introduce implicit assumptions about the behavior.

The three refinements to the reinforcement schedule introduced in this report fit into two categories. The first is an application of the central limit theorem and is independent of the other two. Rather than using just a single sample of the steady-state data, multiple independent samples are used to compute an average transfer function in an effort to improve the signal-to-noise ratio. The other two refinements are improvements in the form of the reinforcement schedule used to extract the transfer function. Unlike Palya et al. (1996), we used two or three VI pulses during the transfer function measurement phase, and we delayed the first pulse onset from the trial's start for all phases. Both refinements were driven by the analytic properties of linear analysis and the empirical properties of the schedule behavior. Employing them alleviates a pair of artifacts that make predicting behavior on a new schedule more difficult. Although these two refinements interact to some extent, they can be considered sequentially.

We increased the number of steady-state sessions available for the analysis reported here, particularly for the phase specifically designed for transfer function measurement. Assuming the dynamics were unchanged, the extra data allowed the comparison of three different methods of transfer function computation. All three methods are similar; they are more like variations on a theme than unique estimation algorithms. The three transfer functions, once extracted, are used identically, as shown on the right half of Figure 1, to predict the behavior that should be supported by a second schedule. The first method estimates the transfer function based on a single sample of the steady-state data exactly as shown on the left half of Figure 1. The second method is based on extending the number of steady-state sessions. In this second variation, five independent samples of the steady-state data are used to compute five distinct transfer functions. Each of the five transfer functions is extracted as shown on the left half of Figure 1. The five transfer functions are then averaged in the frequency

domain to yield a mean transfer function estimate. The third method is similar to the second in that it also combines more than one transfer function to form an average. The difference is that we now estimate transfer functions based on data from two different experimental conditions and then average in the frequency domain. For this variation to be successful, the experiment must contain at least two experimental conditions that are suitable for transfer function extraction, as well as a third experimental condition to serve as a test. As a side benefit, the greater number of steady-state sessions required by the second method also provides an opportunity to estimate the precision of the repeated trials local average without requiring prior assumptions about the forms of the distributions.

Palya et al. (1996) showed that determining a transfer function on trials containing a single reinforcement pulse is problematic because it results in reinforcement (input) components of zero amplitude at behaviorally important low frequencies. The zero amplitudes occur at frequencies corresponding to the reciprocals of the pulse duration and at all higher harmonic frequencies. Because there was no information for those frequencies, the predictive accuracy of the procedure was reduced. Palya et al. pointed out that if the procedure filled those zero amplitude frequencies, the resulting transfer function would make better predictions. This artifact's linkage to the reciprocal of the pulse period suggested its solution. Kessel (1998) developed the mathematical expressions that describe how the addition of a second pulse can fill these ''amplitude holes.'' The core attribute required of the second pulse is that its duration differ from the first. Frequencies not present in one pulse are filled in by the other. This refinement improves the transfer function estimation by using reinforcement schedules with more uniform spectral content across the low frequencies. The more a schedule yields flatter spectra, the lower the likelihood of divide-by-zero artifacts. The interval between the two pulses also affects the low-frequency component amplitudes, although less dramatically. We would not, however, claim that the schedules developed for this report are the best possible schedules for the measurement of a transfer function. Rather, the set of schedule parameters proved adequate for the problem at hand.

In Palya et al. (1996), there were usually five 1,000-s trials per daily session with the intertrial boundary marked by a 10-s blackout. The VI pulse always began at the start of a trial. Hence, the blackout marked not only the intertrial boundary but also certain entry into VI 20 s. Because of this unintended signal, the rates changed instantaneously from zero during the blackout to a high rate during the VI. This produced an artificially sharp onset transition in the response wave form. In the present experiment a second or third rising edge could be present later in the trial. This situation would lead to a confounding effect in measuring the onset dynamics in which data from a signaled edge (after exit from the blackout) are mixed with data from an unsignaled edge (at an arbitrary time in the trial). There would be difficulties when trying to predict the effect of purely unsignaled onset edges. We removed this problem by inserting a short extinction period ahead of the first VI pulse. As a result, response rate at the beginning of the trial started at a relatively low level and exhibited an orderly increase up to the time the unsignaled VI went into effect. This initial low rate also prevented a discrepancy between the rates at the onset and end of the Fast Fourier Transform (FFT) sampling window. The results of Palya et al. were unaffected by this confounding effect because they studied only variations in the falling edge timing and the magnitude of a single pulse.

METHOD

Subjects

Five adult experimentally naive pigeons obtained from a local supplier were used. They were housed under continuous illumination in individual cages with free access to water. All were maintained at approximately 80% of their free-feeding weights by limited feeding with pelletized laying mash.

Apparatus

Five experimental chambers were used. The interior of each was a box (30 cm by 30 cm by 34 cm). An unfinished aluminum panel served as one wall of the chamber; the other sides were painted white. The aluminum

panel had a feeder aperture 5 cm in diameter, medially located 10 cm above the grid floor. Three response keys, 2 cm in diameter, were located 9 cm apart, 29 cm above the grid floor. Only the center key was used. It required approximately 0.15 N to operate. The key was transilluminated green by a stimulus projector containing the Rosco 86 yellow-green theatrical gel throughout all phases of the experiment, with the exception of the blackouts and reinforcement. Two houselights were located on the stimulus panel 32 cm above the grid floor. The lamps were shielded such that their light was directed towards the ceiling. Ventilation was provided by an exhaust fan mounted on the outside of the chamber. A white noise generator provided ambient masking noise.

Stimulus events were controlled and key pecks were recorded by a computer system (Palya & Walter, 1993). The computer archived the time of each stimulus and response event in 1-ms ''ticks.'' Subsequent data extraction and analysis routines provided the derived data. Complete raw data event logs of all research are maintained for 10 years and are available from the authors.

Procedure

Each pigeon was trained to approach and eat from the food magazine within 3 s on three consecutive presentations. They were then autoshaped to peck an illuminated key for five to six sessions. Once stable responding had been established, the experimental phases began. Throughout the course of the experiment, each daily session typically contained 45 to 55 food presentations. The precise number of food presentations, however, was determined by the pigeon's body weight that day and was adjusted to be just sufficient to maintain the pigeon at its 80% weight. Pigeons that required more food to maintain their weight were postfed no less than 60 min after their daily sessions. Each phase continued until two criteria were satisfied in sequence. First, the pigeon's behavior had to reach a steady state in which response rate showed no apparent session-to-session trends, as judged by visually inspecting average response rate plotted as a function of sessions. After this stability criterion had been met, the phase then continued until the number of sessions required by the various averaging techniques had been obtained and the change could be fitted within the constraints of other laboratory activities.

Following exposure to initial training conditions, the pigeons were exposed to the four phases of the experiment. All four used repeated exposures to 2,000-s trials separated by 10-s blackouts. Within each trial, circumscribed but unsignaled periods of reinforcement availability (pulses) occurred, during which a VI 20-s schedule operated. Outside the VI pulse durations, pecks had no effect (an extinction schedule). The VI schedules were composed of exponential distributions of interreinforcer intervals (IRI) (Fleshler & Hoffman, 1962). The trial procedures constituted mixed reinforcement schedules. There were typically three trials in each daily session.

The modulation patterns that define the unsignaled VI pulses within a trial are shown schematically for each pigeon and phase against the 2,000-s trial length in Figure 2. The modulation patterns do not contain the local variations in the reinforcement rate that are present during the VI pulses. The local variations will have to be incorporated into the analysis because they give the reinforcement input a significant distribution width. For descriptive purposes, however, it is convenient to regard the simplified patterns in Figure 2 as the input signals. The onset times and completion times for the VI pulses shown in Figure 2 are given in Table 1. As noted above, Fourier's theorem requires that the trials contain a range of input frequencies determined by the pulse widths and spacings selected for each phase. Of the experiment's four phases, only Phase 3 was expressly intended for the computation of a transfer function as shown on the left side of Figure 1. We used a two-pulse procedure in the measurement of the transfer function for all pigeons except Bird 568. Bird 568 was on a three-pulse procedure during Phase 3. Phase 3 was in effect for a relatively longer period to provide enough sessions to make five independent estimates of the transfer function. Phases 4 and 5 served primarily as the test phases shown on the right side of Figure 1. During Phases 4 and 5, all pigeons were exposed to a procedure containing either three long pulses closely spaced or three short pulses widely spaced (counterbalanced across pi-

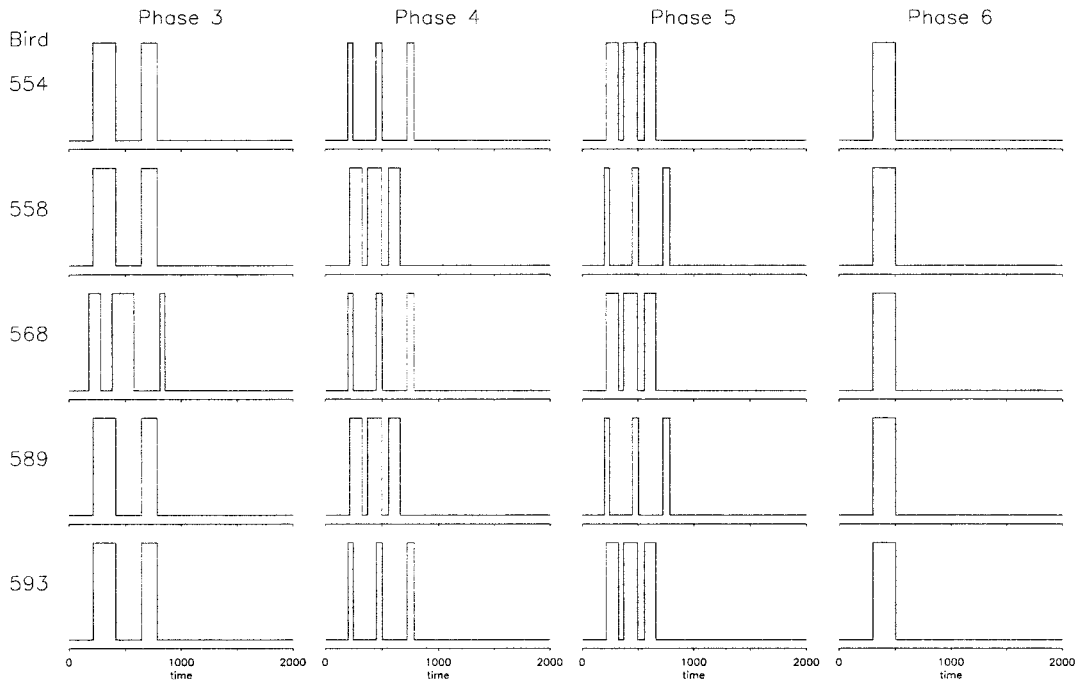

Fig. 2. A schematic of the reinforcement profile for each pigeon and phase of this experiment. These reinforcement profiles result from the combination of VI burst and intervening extinction. All trials have a common duration of T_{trial} = 2,000 s and a 10-s blackout between trials. The left column shows the profiles used exclusively for the measurement of the pigeon's transfer functions. The other three columns show the three-pulse and single-pulse phases used to test the prediction of a linear analysis. The local variations present in the actual reinforcement rate have been suppressed for clarity.

geons). Inspection of the behavior suggested that there was no order effect; the behavior supported by the two- or three-pulse procedure was a property of the width and spacing of the pulses and not the presentation order. Although not specifically designed for the purpose, the phases with three narrow pulses widely spaced could also be used to extract a useable transfer function. Phase 6 was a single-pulse schedule common to all 5 pigeons and was used solely as a test phase. The single pulse of Phase 6 was moved 100 s later into the trial to address two minor questions. First, is there a commensurate change in the pigeons' behavior? Second, how well does a linear analysis handle this form of generalization? In both cases, the results were as expected: Both the behavior and the prediction did follow the reinforcement.

RESULTS AND DISCUSSION

In a broad sense, the behavioral dynamics supported by multiple unsignaled VI pulses

Table 1

Start and stop times for pulses during the 2,000-s trials.

Phase	Bird	Start	Stop	Start	Stop	Start	Stop
3	555, 558, 589, 593	211.0	411.0	642.0	783.0		
	568	174.0	279.0	381.0	574.0	807.0	854.0
4	554, 568, 593	198.2	245.2	448.2	503.0	721.6	784.2
	558, 589	213.8	323.2	370.1	495.2	557.7	659.2
5	554, 568, 593	213.8	323.2	370.1	495.2	557.7	659.2
	558, 589	198.2	245.2	448.2	503.0	721.6	784.2
6	All	299.8	503.0				

intermixed with periods of extinction have parallels to the behavioral dynamics reported by Palya et al. (1996). There are, however, some new features. To highlight the similarities and differences, we will first consider the results from a phenomenological viewpoint, and then discuss the performance and limitations of the linear analysis. Similarly, although the data-reduction methods follow Palya et al. (1996), some technical aspects of their use are specific to this experiment. This material can be found in the Appendix.

The steady-state behavior showed orderly responding with a correlation to the VI pulses for all pigeons during all four phases (Figure 2). The top panel of Figure 3 shows the individual response rates in each time bin for all trials over the final 20 sessions of Phase 4 for Bird 558 as a dot plot (Palya, 1991, 1992). The bottom panel is the corresponding plot for Bird 558 during the final 20 sessions of Phase 5. The repeated trials local average is indicated with the solid line, and the pulse intervals are indicated by the dark gray bars. In much the same way as the rectangular modulation profiles used in Figure 2, our use of dark gray bars simplifies the depiction of reinforcement. A dot plot of the reinforcers would show a constant rate of zero outside the VI pulse interval and a rectangular band centered at the mean VI during the interval. The light gray bars indicate the duration of the maximum IRI added to the end of the scheduled pulse. This ''+1 max IRI'' region is terminated with a dashed line.

Figure 3 shows data typical of the dynamic behavior supported by the VI pulse schedules. The response rates increased during the initial extinction between the intertrial blackout and the first period containing reinforcers. Response-rate decrements typically occurred at the start of the interpulse extinction periods, followed by an anticipatory increase in rate up to the next period containing reinforcers. In general, the drop in average response rate during an interpulse extinction increased as the duration of extinction increased. Response rates within the long final extinction typically fell to a minimum relatively promptly after the VI pulse completion and then subsequently increased as the next blackout neared. Compared to the steady, featureless cumulative records traditionally obtained for a continuous VI schedule, a VI pulse results in substantially more local variation in behavior during the pulse. For some multipulse phases, the final VI pulse was also followed by a ''ring'' or ''echo'' of temporally localized enhanced responding within the long final extinction. Although the broad form of the supported behavior in all 5 pigeons was similar, the responding of each pigeon was distinctive, in agreement with Palya et al.'s (1996) results. Bird 568 was sensitive to even the shorter extinctions, whereas Bird 593's responding tended to show less variation even at the wider pulse separations. Considering the behavior of all 5 pigeons as a group, it appears that the underlying process governing the behavioral dynamics we observed is the same as seen in the results of Higa (1996), Higa and Pierson (1998), and Horner, Staddon, and Lozano (1997). In all four studies a sharp change in local reinforcement rate produced a change in response rate with roughly a 30-s rise (or fall) time. (See the discussion following Equations 4 and 5 in the Appendix for the effect a 30-s transition time had on the experimental design.)

All 5 pigeons maintained surprisingly high response rates during the long final extinction period despite hundreds of sessions of exposure to the complete lack of reinforcers in the final approximately 1,200 s of the trial. Unlike Palya et al.'s (1996) results, even the single pulse procedure used in Phase 6 of the present study supported some responding throughout the final extinction period. The responding during the long final extinction was quite variable between individual trials. It is unclear whether the appearance of this responding following a single pulse was due to adding an extinction interval between the trial boundary blackout and the first unsignaled reinforcement pulse onset, was an order effect caused by the prior multiple pulse phases, or was the result of some other behavioral process. As can be seen in Figure 3, the lowest rate occurred shortly after the last VI pulse (or the ring) of the trial, rather than at the end of the long extinction period. This finding was inconsistent with behavioral momentum or reflex-reserve-like notions, which would suggest that the strengthening effect of the reinforcement periods dissipates with increasing exposure to extinction. Neither was it consistent with stimulus control notions, suggesting that it should be increasingly ap-

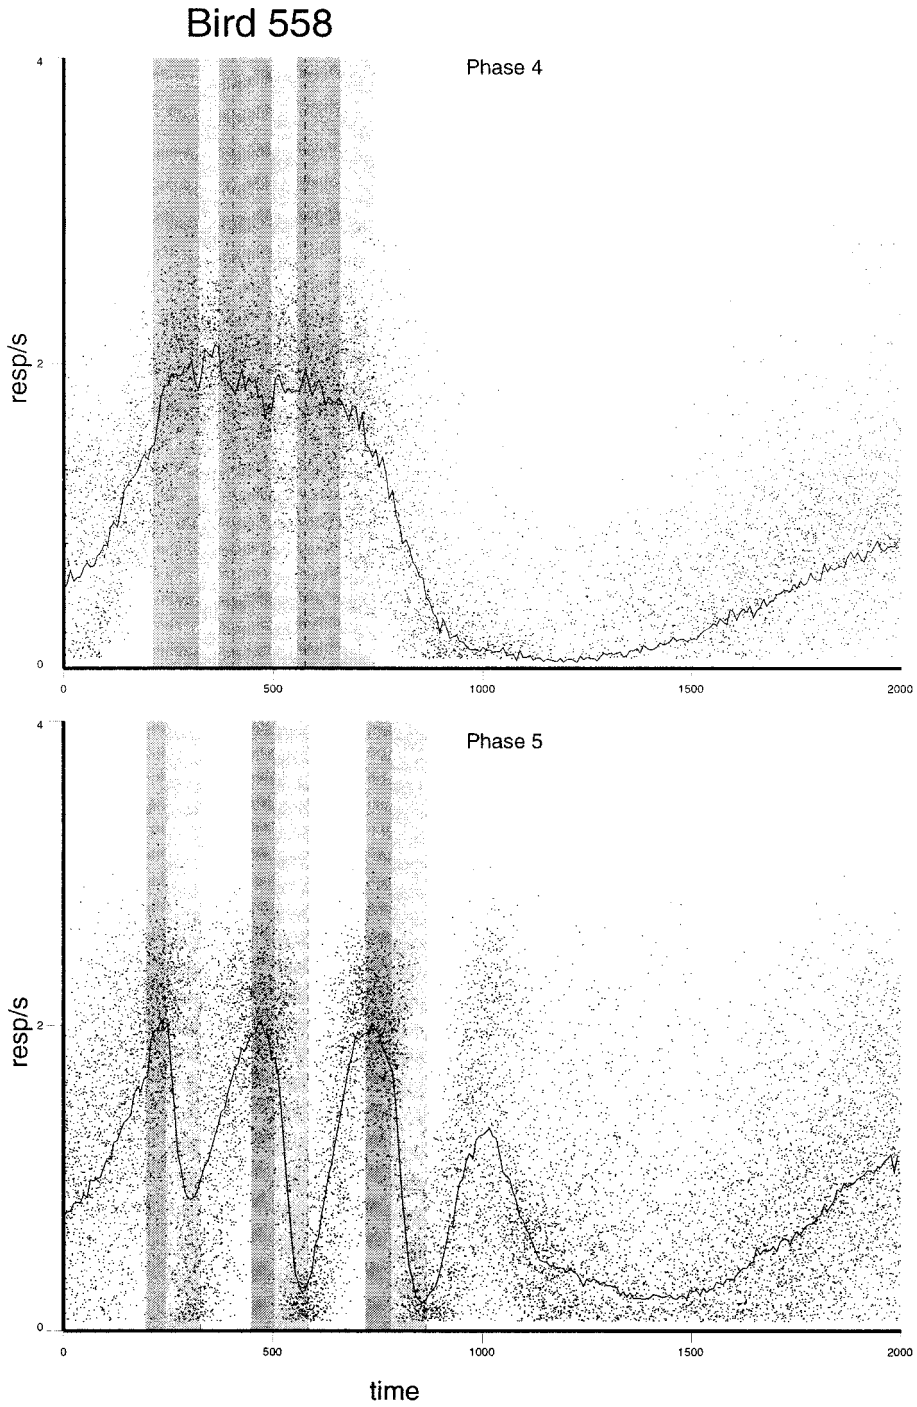

Fig. 3. Dot plots showing the responding of Bird 558 under the narrowly spaced and widely spaced procedures of Phases 4 and 5. The measured response rate for each of the 256 time bins for each trial during each of the last 20 sessions is plotted as a dot. The mean response rate within each bin is the repeated trials local average and is depicted with a solid line. The VI pulses containing food availability are designated with dark gray bars. Light gray bars show the length of the maximum IRI and are appended to the reinforcer availability period. The "+1 max IRI" region is terminated with a dashed line.

parent that no further reinforcers would be forthcoming as the final extinction period elapsed. A 1,200-s period of extinction is unlike a VI 20-s schedule. The observed behavior is consistent with Palya's (1993) bipolar model of behavior across temporal intervals. His model predicts high rates at the point just before the reinforcer and predicts the lowest rates at the point most removed from the upcoming reinforcer. In the present procedure, the termination of the last reinforcer pulse would be predicted to control the lowest rate in the trial.

A notable effect generated by some, but not all, procedures is illustrated in the bottom frame of Figure 3. When the VI pulses were widely spaced, a localized increase in responding (a ring) occurred approximately 200 s after the final pulse. The ring's rising edge appears in form to be anticipatory responding before a nonexistent fourth pulse. By contrast, the top frame of Figure 3 illustrates the lack of a ring typical for procedures with narrow gaps between the VI periods. Providing a satisfactory description of the ring may well represent a challenge beyond the capabilities of a linear analysis and probably for many other approaches as well. We will consider this effect and its implications further once the quantitative performance of the linear analysis predictions is available for context.

Each of the 256 time bins in Figure 3 contains a distribution of response rates across trials. The form of the distribution varied as a function of time throughout the trial. The distribution of response rates is broad and centrally peaked about the repeated trials local average during the VI pulses. During the extinctions, the response-rate distribution is single sided with a skew towards zero. In neither case are the response-rate distributions particularly well approximated by a Gaussian distribution. Given the form of the distribution, how precise are the repeated trials local averages, and how quickly does adding additional trials improve its precision? The five independent samples of behavior measured in Phase 3 can provide an answer. For each of the 256 bins, we determined the range between the minimum and maximum of Bird 558's five repeated trials local average rates. This minimum-to-maximum range of the five rates is narrower than that of the overall distribution width seen in Figure 3. One can then compare one half of this directly measured minimum-to-maximum range with that expected from the standard error of the mean, $SEM = SD/\sqrt{N}$, where *SD* is the standard deviation for Bird 558's overall distribution. When averaged over all the 256 bins, the *SEM* is a reasonable approximation to the directly measured range, assuming that N is the number of sessions ($N = 20$) and not the number of trials. Extending this approximation, one can use the individual-bin *SEM* values squared for the $\sigma^2_{B(t_i)}$ term in Equation 8 (see the Appendix). Although the *SEM* slightly underestimates the average width, it is a convenient approximation because it is determined by the data from a single sample. More locally within the trial interval, the measured range has a noticeably different time dependence. The extinction and the back edges of the response pulses show greater variability than expected from the *SEM*.

Although the linear analysis to follow is focused on the averages of the distributions (i.e., the solid line in Figure 3) and, to a lesser extent, their widths as measured via the variances, there is more to the data. As noted, the repeated trials local average and standard deviation do not fully capture the variations in the form of the distributions as a function of time throughout the trial. These variations in the steady-state form of the distribution can be more naturally shown with a color three-dimensional projection plot,[1] a series of profile plots (a waterfall plot), or a color spectrogram plot rather than Figure 3, which highlights the average character. A second important characteristic of behavioral dynamics that is eliminated in constructing a repeated trials local average is the trial-to-trial sequential structure. This molecular structure in behavior is best shown by plotting all trials individually and playing them as a movie.[2]

Figures 4 through 8 show the linear analysis of the results as well as the response and reinforcement rates. Each figure shows measured and predicted response rates across the

[1] A sample three-dimensional color projection plot is available at http://www.jsu.edu/depart/psychology/sebac/linear-modeling/. The equivalents to Figure 3 for all pigeons and all phases are also available at this URL.

[2] A set of QuickTime movies for all pigeons and all phases are available at http://www.jsu.edu/depart/psychology/sebac/linear-modeling/.

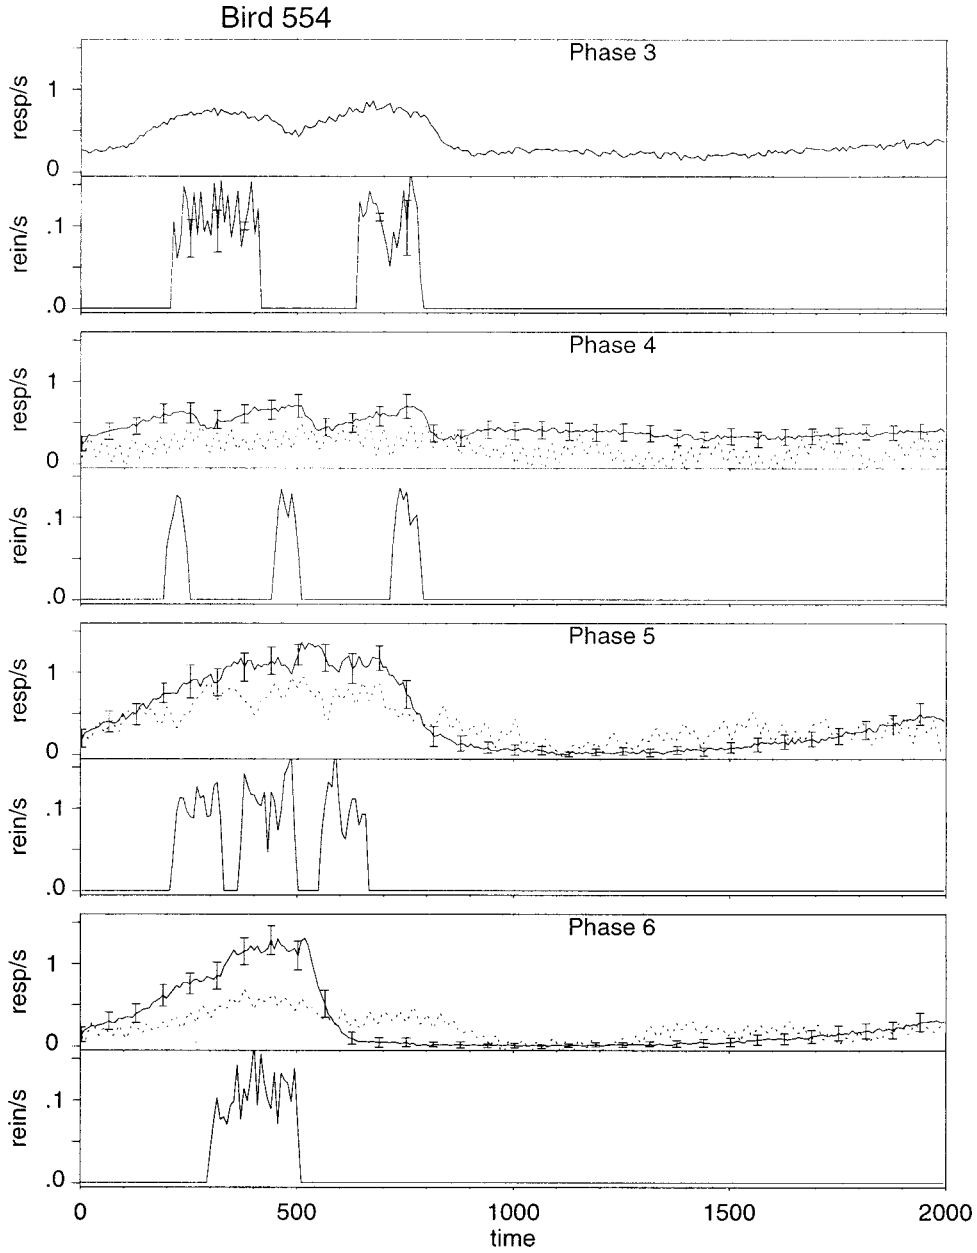

Fig. 4. The obtained and predicted responding across a trial for Bird 554 for each phase of the experiment. The upper frame in each pair depicts the obtained response rates with a solid line and the prediction made by a linear analysis with a dotted line. The *SEM* of every 8th time bin's measured response rate is shown by the superposed error brackets for Phases 4, 5, and 6. The remaining error brackets, not shown for clarity, were the same local magnitude as those shown. The lower frame in each pair of frames indicates the obtained reinforcement rate in each time bin of the interval. The minimum to maximum range of the repeat trials local average reinforcement rate from the five samples of Phase 3 is superposed as an error bracket on every 8th time bin. The predictions all used the single-sample transfer function b554_t_fun_p3_v1.dat.

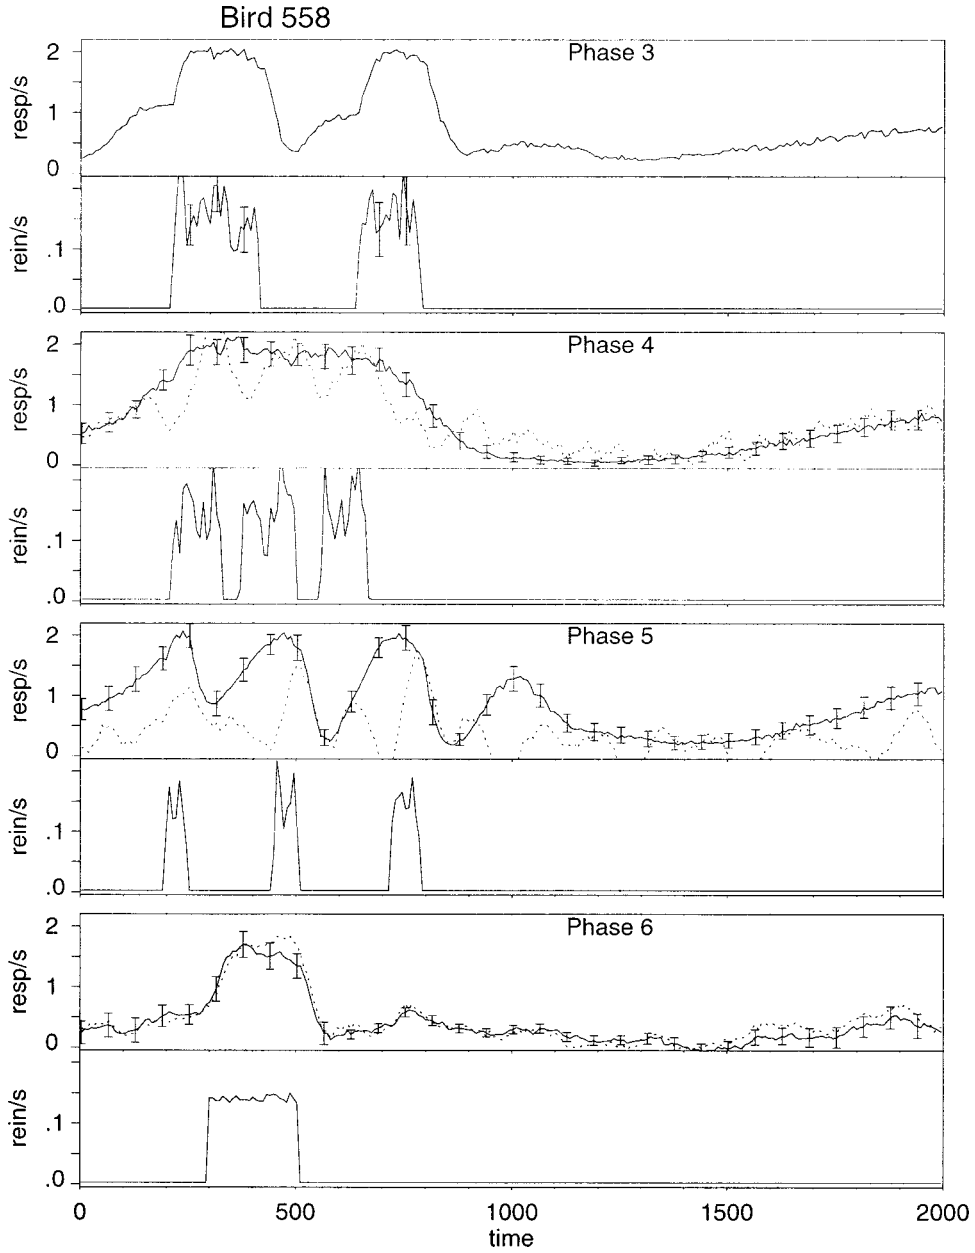

Fig. 5. The obtained and predicted responding across a trial for Bird 558 for each phase of the experiment. The predictions all used the single-sample transfer function b558_t_fun_p3_v1.dat. Details as in Figure 4.

four phases. Paired with these response-rate plots are panels showing the measured reinforcement rates. The measured rates are repeated trials local averages computed from all trials during the final 20 sessions of each phase. Excluding a special-case self-prediction for Bird 558 used to establish the propagated uncertainty due to the variance of the reinforcement-rate distribution, Phase 3 served as the basis for transfer function estimation and no predictions are shown. Once the transfer function had been obtained from Phase 3, predictions for average behavior during Phases 4, 5, and 6 could be made, as shown in Figure 1. The linear analysis predictions are shown for Phases 4, 5, and 6 with

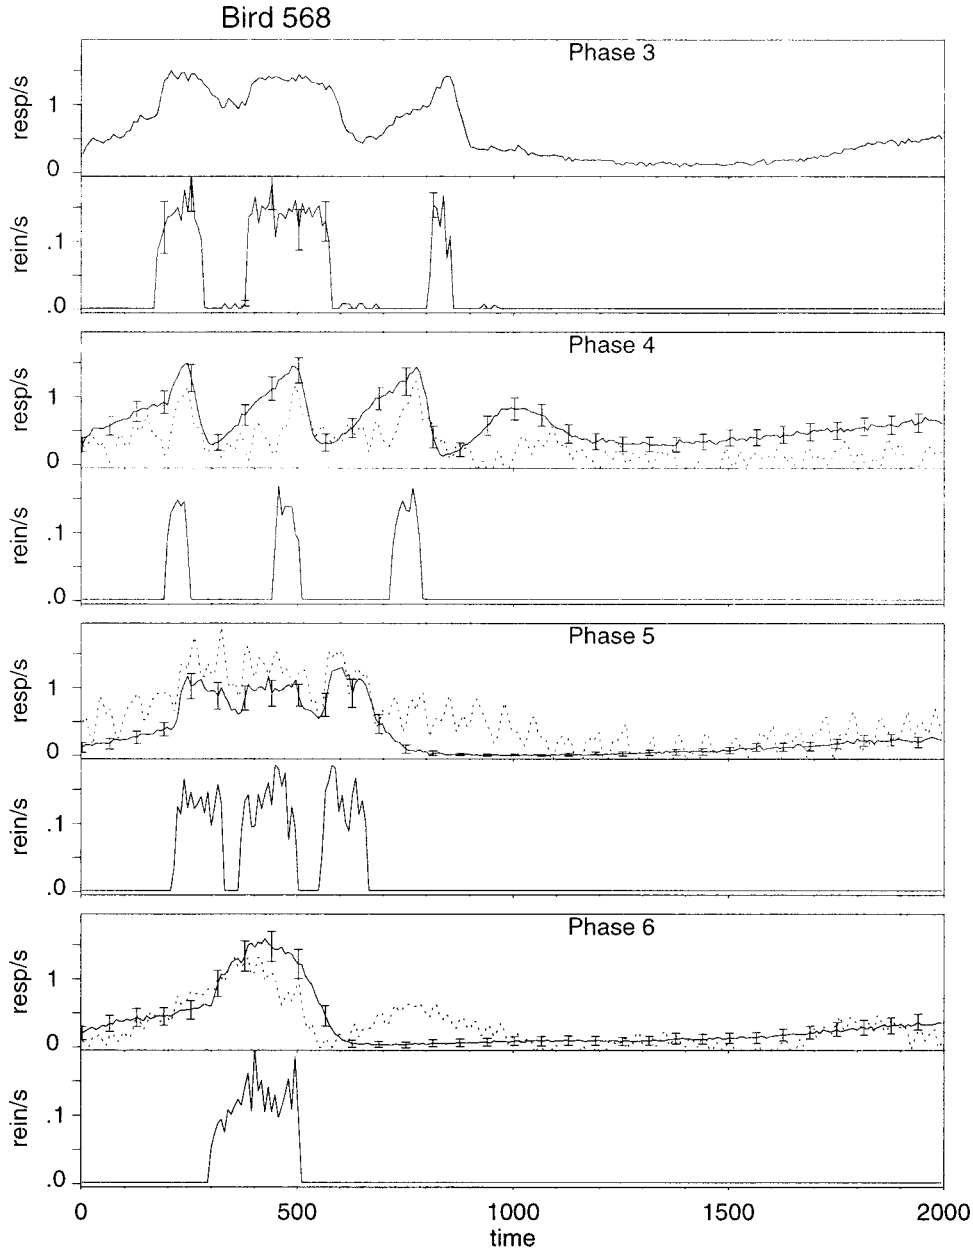

Fig. 6. The obtained and predicted responding across a trial for Bird 568 for each phase of the experiment. The predictions all used the single-sample transfer function b568_t_fun_p3_v1.dat. Details as in Figure 4.

dotted lines. Note that the transfer functions used for the predictions are denoted by their data file names (e.g., b554_t_fun_p3_v1.dat contains Bird 554's transfer function as estimated from Phase 3, Single Sample 1). Each of the three predictions in Figures 4 through 9 is made from a single sample of the pigeon's behavioral dynamics measured over the last 20 sessions of Phase 3. Bird 554 had both the lowest characteristic response rate and the softest ring of the 5 pigeons. The linear analysis, of course, is unchanged by the presence of the subject-to-subject variations in responding, such as those apparent in Figures 4 through 8. It merely asserts that each pigeon's behavioral dynamics measured in

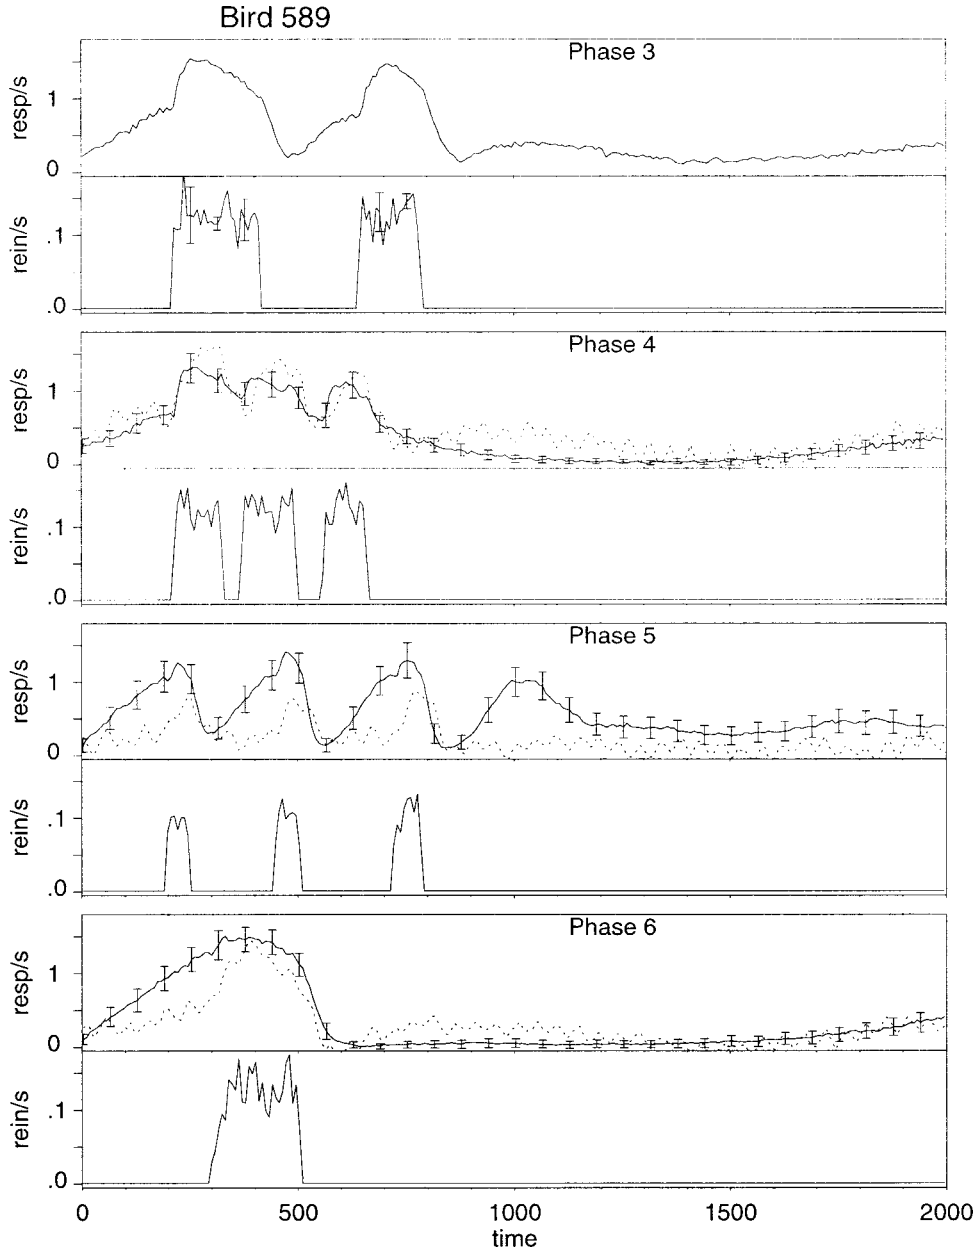

Fig. 7. The obtained and predicted responding across a trial for Bird 589 for each phase of the experiment. The predictions all used the single-sample transfer function b589_t_fun_p3_v1.dat. Details as in Figure 4.

Phase 3 can be extrapolated to Phases 4, 5, and 6 by the predefined algorithm.

Figures 4 through 8 also show uncertainty brackets in four of the eight panels. The vertical bars show the *SEM* of the measured response rates for every eighth of the 256 time bins in Phases 4, 5, and 6. In general, the *SEM*s for the long final extinction are comparatively smaller than those during the pulses. The corresponding smaller variances will increase the effect on the χ^2_ν test of response-rate discrepancies during the long extinction interval. The vertical bars on the Phase 3 measured reinforcement rates are taken directly from the observed distribution width of the five repeated trials local averages, rather

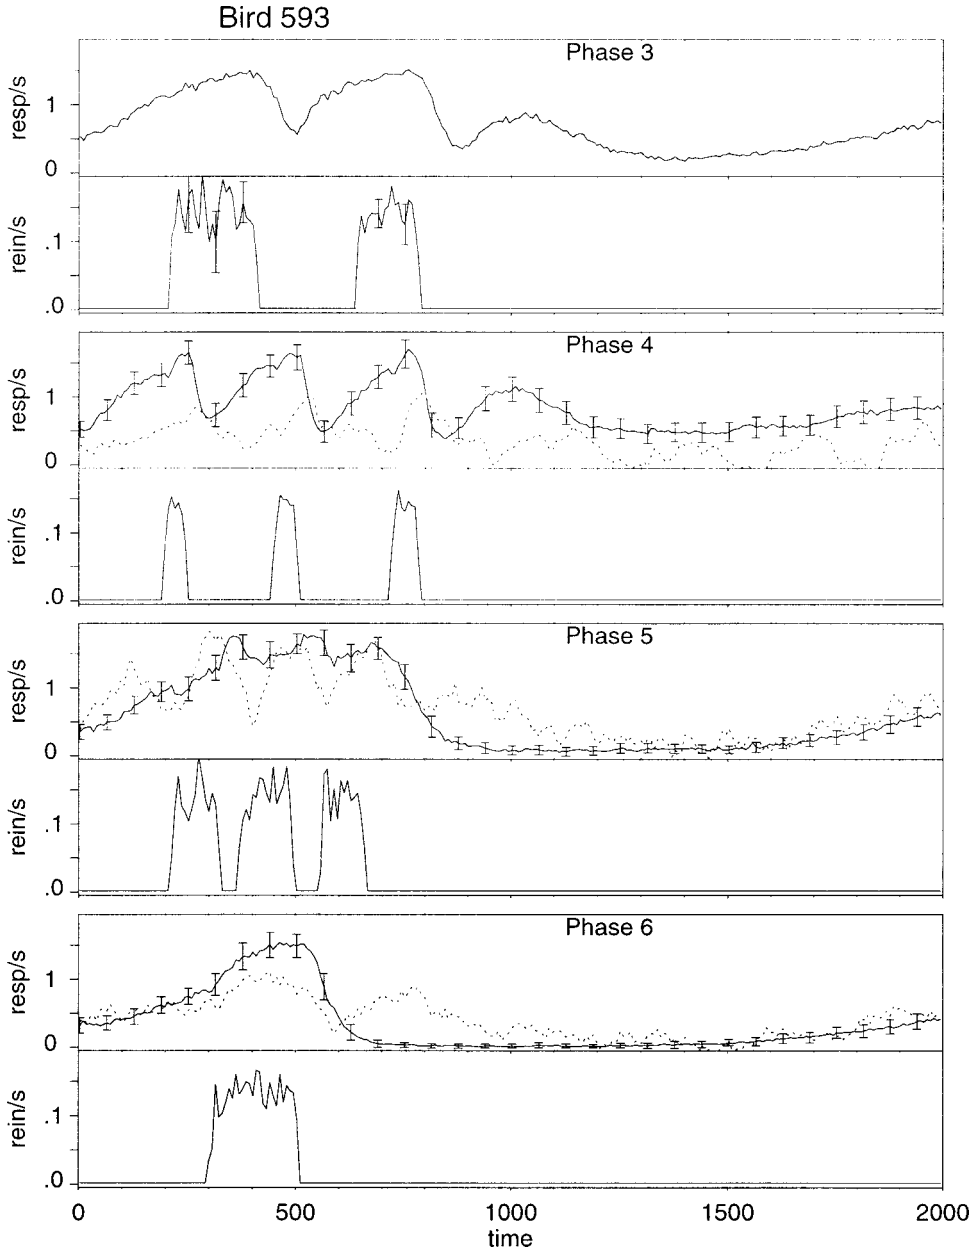

Fig. 8. The obtained and predicted responding across a trial for Bird 593 for each phase of the experiment. The predictions all used the single-sample transfer function b593_t_fun_p3_v1.dat. Details as in Figure 4.

than an *SEM* estimate. These reinforcement rate uncertainty brackets are used to estimate the $\sigma^2_{\bar{B}^{\mathrm{pred}}(t_i)}$ term in Equation 8 (see the Appendix).

Of the three predictions for the responding of Bird 554 in Figure 4, Phase 5 is the best, with Phases 4 and 6 lower in fidelity. All three predictions are faithful to the basic form and magnitude of the response rates. Based on the single Phase 3 sample of Bird 554's behavioral dynamics, the linear analysis correctly predicts lower and relatively less differentiated rates in Phase 4 compared to Phases 5 and 6. None is perfect, particularly in the fine structure. A rapid, approximately 32-mHz, quasiperiodic oscillation appears in

the predicted response rate that is not seen in the measured rates. Those fluctuations that do exist in the measured response rates are more localized, spiky, and noise-like in character. Many of the same observations about the quality of the predictions for Bird 554 also apply to the other 4 pigeons. Clearly, a linear analysis that estimates the transfer function by the simple division shown in Figure 1 is incomplete. For about half the predictions, sharp resonances in the transfer functions result in noticeable quasiperiodic oscillation. There was also a modest tendency for the predicted response rates to fall below the observed rates in all 5 pigeons. Still, using a schedule such as that in Phase 3, that is specifically tuned for extracting a transfer function has improved the predictions. Compared to the predictions of Palya et al. (1996), these predictions have fewer excursions into behaviorally impossible negative response rates. Comparing the measured and predicted plots provides a useful qualitative sense of the predictions; to develop a quantitative assessment of these predictions or to quantitatively compare two predictions requires a χ^2_ν test.

As discussed in the Appendix, using χ^2_ν entails a dependence on the variance of the residuals between the measured and predicted response rates. There are two aspects of this dependence of particular importance for this report. First, a large discrepancy between obtained and predicted values in a bin with substantial variability does not contribute to χ^2_ν as heavily as a large discrepancy in a bin with little variability. Consequently, prediction discrepancies during the final long extinction typically have a proportionately larger impact on χ^2_ν because of the lower variance that usually occurred during this extinction. As a counterexample, the relatively larger variance brackets for Bird 554 during the long final extinction of Phase 4 make a low χ^2_ν value easier to obtain relative to Phases 5 and 6. Second, the variance may be estimated in a variety of ways, particularly because the distributions depart significantly from Gaussian. We employed two variance estimates in this report.

Palya et al. (1996) used the obtained variance of the response-rate distribution. With reuse of this estimate, we obtain χ^2_ν values with identical significance for both studies, allowing a relative χ^2_ν comparison. As an alternative, we also used a tighter variance estimated from the five Phase 3 repeated trials local average reinforcement and response rates. As has already been noted, the range in the Phase 3 repeated trials local average response rates can be approximated by the overall distribution's *SEM*. The *SEM*, by itself, ignores the contribution to the variances of the residuals from the width of the averaged reinforcement distribution (second term in Equation 8). To estimate the effect of the reinforcement distribution width, we computed a pair of predictions for Bird 558's Phase 3 response rates: one from the minimum and one from the maximum reinforcement rate in each time bin. The spread between the minimum and maximum prediction values is the effect of the reinforcement distribution width propagated through the linear analysis. One half the propagated width, assuming a rectangular distribution, provides an estimate of the prediction width $\sigma_{B^{\mathrm{pred}}(t_i)}$. Combining this prediction precision with the response-rate precision estimated from the repeated trials local average range (the *SEM*) yields an estimate for the total variance via Equation 8. When compared to the standard deviation of the overall response-rate distribution (Phase 3 equivalent of an individual panel in Figure 3), total precision for the residuals is smaller by roughly a factor of $\sqrt{2}$. Consequently, we can now use $SD^2/2$ as a second tighter variance for weighting the residuals in χ^2_ν. This second variance estimate still suppresses any variance attributable to the transfer function estimation. Had some other variance estimate been used, the absolute values of the χ^2_ν would be shifted, although the relative ratios between all values would remain unchanged.

Table 2 gives the χ^2_ν values for predictions based on three different estimates of the linear transfer function: a single sample from Phase 3, a five-sample average from Phase 3, and a two-sample average combining Phase 3 with Phases 4 or 5. Also given are values of the probability q that are associated with the χ^2_ν values, assuming the two variance weightings. The q_{SD} values are computed with the same variance as in Palya et al. (1996). The $q_{SD\sqrt{2}}$ values use the tighter variance $SD^2/2$ weighting discussed above. For either variance weighting, q is the conditional probability that a discrepancy between the measured

Table 2

Prediction χ^2_ν values and probabilities for exceeding chance.

	Bird				
	554	558	568	589	593
χ^2_ν for predicted based on [KESSEL.VI_EXTINC]b5**_t_fun_p3_v1.dat (last sample of Phase 3 only)					
Phase 4	0.334848	0.312560	0.392401	0.648345	0.627654
q_{SD}	1.000	1.000	1.000	0.9999972	0.9999994
$q_{SD/\sqrt{2}}$	0.999988	1.000	0.995	0.00094	0.0034
Phase 5	0.339258	0.787471	6.977813	0.306912	0.632041
q_{SD}	1.000	0.9949	0.000	1.000	0.9999992
$q_{SD/\sqrt{2}}$	0.999979	0.000	0.000	1.000	0.0027
Phase 6	0.651850	0.401629	0.717874	0.393893	2.273314
q_{SD}	0.9999964	1.000	0.99979	1.000	2.0×10^{-27}
$q_{SD/\sqrt{2}}$	0.00078	0.991	5.8×10^{-6}	0.995	0.000
χ^2_ν for predict based on [KESSEL.VI_EXTINC]6b**_t_fun_p3ave.dat (five Phase 3 samples averaged in frequency)					
Phase 4	0.270221	0.429029	0.359487	0.669350	0.662373
q_{SD}	1.000	1.000	1.000	0.99998	0.999992
$q_{SD/\sqrt{2}}$	1.000	0.95	0.9998	0.00024	0.00038
Phase 5	0.241188	0.766648	5.023332	0.298419	0.595673
q_{SD}	1.000	0.998	0.000	1.000	0.9999999
$q_{SD/\sqrt{2}}$	1.000	6.0×10^{-8}	0.000	1.000	0.019
Phase 6	0.630413	0.413695	0.559845	0.418702	1.884797
q_{SD}	0.9999993	1.000	1.000	1.000	4.4×10^{-16}
$q_{SD/\sqrt{2}}$	0.0029	0.98	0.091	0.97	0.000
χ^2_ν for predict based on [KESSEL.VI_EXTINC]b5**_t_fun_p3*ave.dat (average of last sample of Phase 3 and Phase 4/5)					
Phase 6	1.793702	1.665171	1.772257	1.952545	6.602450
q_{SD}	9.9×10^{-14}	1.2×10^{-10}	3.4×10^{-13}	6.4×10^{-16}	0.000
$q_{SD/\sqrt{2}}$	0.000	0.000	0.000	0.000	0.000

Note. All predictions are parameter free and have 256 degrees of freedom. Entries of 1.000 or 0.000 denote cases in which differences between limits of the probability range and the actual probability are smaller than can be represented in single precision computer variables. The q_{SD} were computed directly from the listed χ^2_ν using Equation 9. the $q_{SD/\sqrt{2}}$ were also computed using Equation 9, but after first doubling the listed χ^2_ν.

and predicted response rates that small exceeds a chance agreement. Note that $p = 1 - q$ if significance levels are preferred. As defined by Equation 7 in the Appendix, the χ^2_ν values in Table 2 compare the square of the discrepancy between the predicted value and the obtained behavior in each bin in terms of the bin's variance. The smaller the χ^2_ν value, the greater the fidelity of the predicted rate to the obtained rates. Comparison of the single-sample transfer function predictions in the top section of Table 2 to Palya et al.'s (1996) Table 1 shows a marked improvement. As could be expected from the qualitative improvements discussed above, using the Phase 3 schedule for extracting transfer functions has improved the linear predictions from a quantitative standpoint as well.

Our predictions were for a 256-element ordered time series of response rates with no free parameters. For comparison, a typical single-key Herrnstein hyperbola experiment involves fitting six response rates with two free parameters. For $N = 256$, even modest excursions of the nonexistent echo predicted for Bird 568 in Phase 5 and 593 in Phase 6 fail the significance test. As N increases, the χ^2_ν significance test becomes sharper. The middle ground becomes narrower; either a prediction is well above chance or it fails. At such a large N, q has a plateau structure in which once the agreement between the prediction and the observation closes to within a threshold value, then the probability of the prediction exceeding chance draws very near 1.0. Although room for a better prediction still exists at lower χ^2_ν values, once the plateau is reached little increase in q can occur. Conversely, at such large N, larger residuals will rapidly drive q close to 0.0.

Are better predictions with lower χ^2_ν values possible for the data of this report? Almost

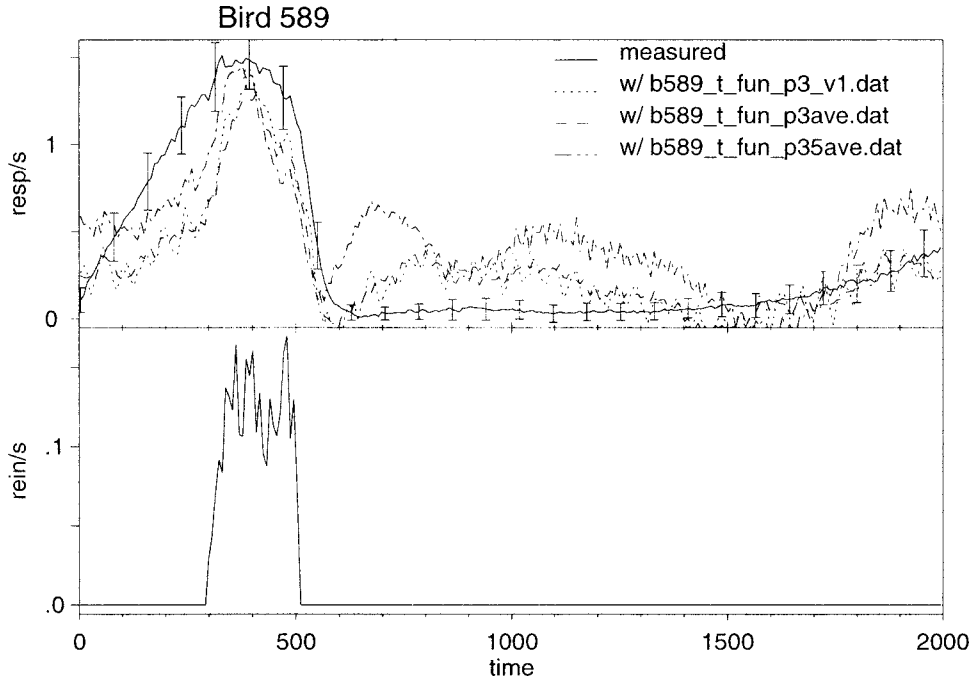

Fig. 9. Phase 6 responding of Bird 589 as measured (solid line) and predicted (dotted and dashed lines as labeled) using three different types of linear transfer function estimate. The single-sample Phase 3 transfer function is b589_t_fun_p3_v1.dat. The five-sample Phase 3 averaged transfer function is b589_t_fun_p3ave.dat. Finally, b589_t_fun_p35ave.dat is an average of two single-sample transfer function estimates, one based on Phase 3 and the second based on Phase 5. The *SEM* of the measured response rate of every 10th time bin is shown by the superposed error brackets. The remaining error brackets, not shown for clarity, were the same local magnitude as those shown.

certainly, although some of the predictions are already good. When some other analysis succeeds in making better predictions, then the quantitative question of how much better will be answered by comparing the χ^2_ν values computed with a common variance estimation. It is important to recognize that our present form of linear analysis yields a partial, but significant, prediction, as opposed to a perfect prediction.

As we have demonstrated, a linear analysis based on a single sample of Phase 3 behavioral dynamics yields predictions that are in significant agreement with the measured behavior. How well does a linear analysis work if it is based on averaging the transfer function estimates from more than one sample of behavioral dynamics? Figure 9 shows the measured behavior for Bird 589 during Phase 6, as well as predictions made with single and multiple averaged transfer function estimates. Similar to Figures 4 through 8, the *SEM* brackets of the measured rates are shown on only every 10th time bin for clarity. Use of an averaged transfer function based on five independent samples of Phase 3 behavioral dynamics (b589_t_fun_p3ave.dat) improved the predictions modestly in terms of qualitative character and from the standpoint of χ^2_ν entries in Table 2. These transfer function estimates have also improved in comparison to both the single-sample estimates and our prior work. Figure 10 shows the five-sample-average transfer function for Bird 589. The top panel gives the amplitude relation between the response rate and the reinforcement rate as a function of frequency. The peak at zero frequency corresponds to the expected rate had the pigeon been on a constant VI 20-s schedule instead of the pulse procedure. The shoulders that slope away from that peak show the decreasing effect (or roll off) of faster rates of change in local reinforcement rate on response rate. The decreases away from the zero-frequency peak are markedly smoother for the five-sample-average transfer function. The phase relation shown in the lower panel shows how local changes in the reinforcement rate are either

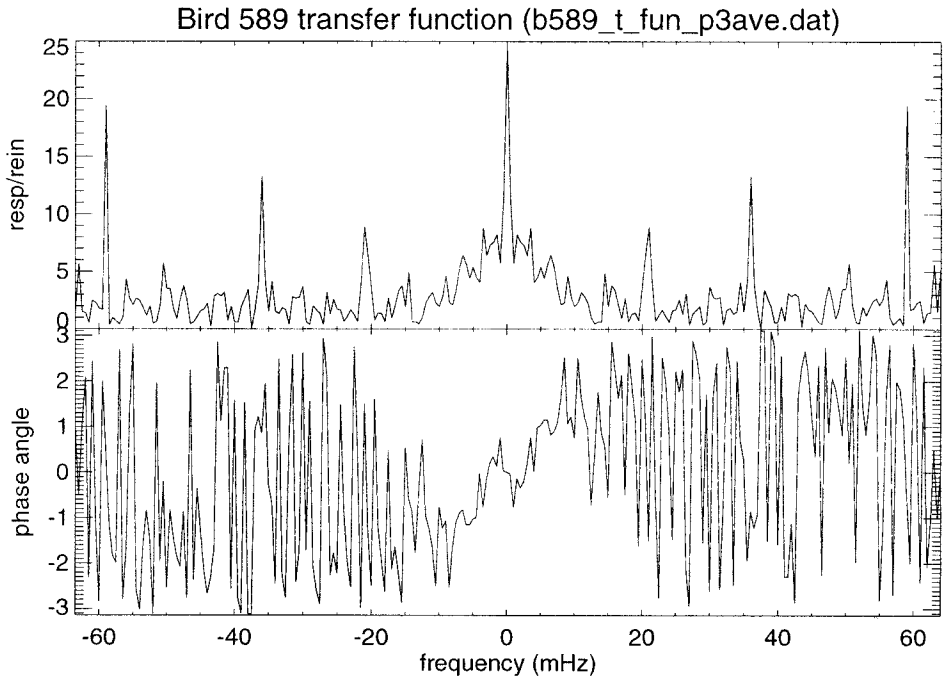

Fig. 10. The magnitude (upper panel) and phase (lower panel) for the five-sample transfer function for Bird 589. The linear transfer function estimate (b589_t_fun_p3ave.dat) is based on five samples of Phase 3 steady-state behavior and the associated reinforcement rates.

anticipated or lagged by the response rate. The five-sample-average transfer function also has a smoother phase relation. The three pairs of narrow spikes in the amplitude plot (upper panel) are likely vestiges of division-by-zero amplitude artifacts and are noticeably smaller than would be seen in the single-sample transfer functions. Consequently, the five-sample prediction suppressed the rapid quasiperiodic oscillations when compared to the single-sample prediction in Figure 9 (b589_t_fun_p3_v1.dat). The suppression of these oscillations for Bird 589 in Phase 6 with a five-sample prediction was also obtained for the other pigeons. A comparison of the top and middle portions of Table 2 indicates that the smoother predictions made by averaging five Phase 3 transfer functions typically, but not invariably, decreased χ^2_ν. The magnitude of the improvement was small, and in 5 of 15 cases the χ^2_ν values of the predictions were worse when using five-sample averaged transfer functions. The second approach, which predicts Phase 6 response rates[3] based on an averaged transfer function that combines a single-sample Phase 3 estimate with a single-sample estimate from either Phase 4 or 5, fails from a qualitative standpoint (b589_t_fun_p35ave.dat in Figure 9) or by using χ^2_ν (bottom section of Table 2). Compared to the other two predictions, including a transfer function estimated from the behavioral dynamics supported by three narrow pulses widely spaced further exacerbates the erroneous prediction of a ring following the single pulse in Phase 6.

The erroneous prediction of a ring in the response rate in Phase 6, and the increase in that error when the three-pulse widely spaced dynamics were included in the predictor, is suggestive of the factors controlling that ring. This potential explanation follows from the

[3] Note that although the Phase 3 reinforcement schedule was specifically designed to extract transfer functions, Phases 4, 5, and 6 were primarily intended as test phases. The phase with three narrow pulses widely spaced is the only one of the three that could be used to extract a transfer function. The other two phases had reinforcement spectrum channels with near-zero amplitudes. As a consequence, we could apply the approach of averaging transfer functions determined during two different phases only to the prediction of Phase 6.

contingency differences between the three procedures involved. As can be seen in the lower frame of Figure 3, in the widely spaced multipulse procedure, the pulses were separated by substantially more than the maximum IRI of the VI schedule. This would result in the regular reinforcement of key pecking following the elapse of a fixed period of extinction. If this were the case, then the ring was a spurious repetition of that response pattern following the third pulse. This interpretation is consistent with the procedure and results illustrated in the upper frame of Figure 3. In that case, the closely spaced pulse procedure had pulses that were separated by less than the longest scheduled IRI in the VI schedules. It could be argued that this was effectively a single pulse and did not support fixed-interval (FI)-like behavior in the interpulse intervals. This, in turn, would result in no ring following the narrow spacing procedure. Following either the last reinforcer in the narrow spaced procedure or the nonreinforced FI-like scallop after the third pulse (the ring) in the widely spaced procedure, another rate drop followed by a rate increase up to the first pulse of reinforcers in the next trial might be expected. By this logic, a linear analysis that averages transfer functions from a widely spaced three-pulse procedure and a two-pulse procedure to get a transfer function would be expected to have a substantial departure from the actual behavior supported by a single-pulse procedure.

We also could have used the second-order Kubo-Bass term (McDowell et al., 1993) to express the discrepancy between the obtained and predicted behavior, instead of contingency changes. The second-order term describes behavior supported by correlation interactions between reinforcement at two different times. Such correlations exist for time lags of approximately the interpulse interval, corresponding to relatively low frequencies, for either two- or three-pulse procedures. It is plausible that the ring is a result of correlation-driven behavior, provided the appropriate interpulse timing occurs. By contrast, a single pulse is isolated and would not be expected to support a well-defined ring via this mechanism. Although the contingency change and the second-order Kubo-Bass term descriptions use different language, the two descriptions are basically equivalent. The behavior-analytic approach posits a contingency change that exceeds the generalization capabilities of a linear model. The second-order Kubo-Bass description explicitly requires a nonlinear dependence of current behavior on past reinforcement history. Both are appeals for the inclusion of some nonlinearity in describing behavioral dynamics. It is worthwhile to note that in physical systems the appearance of such correlation-generated echoes is a standard hallmark of a nonlinear excitation.

CONCLUSIONS

Our approach was to extract and account for sources of variance in a linear analysis of average steady-state behavioral dynamics. Simple linear models can account for a substantial portion of an organism's dynamics without explicit reference to the effects of contingencies of reinforcement or feedback functions. Absent a preferred model, the linear analysis prediction algorithm described here may be the standard by which competitor models of behavioral dynamics, when developed, can be judged. In the short term, there are a number of ways in which other models could outperform the approach of Figure 1. As examples, one could (a) demonstrate predictions better than ours with a reduction in the residual discrepancy (measured by a neutral statistic such as χ^2_ν), (b) accurately describe the complete distribution of behavior illustrated in Figure 3, or (c) operate successfully at higher precision or over a wider generalization range. In the long term, whatever description finally holds for steady-state behavioral dynamics should serve as the limit case for models of the acquisition of behavior.

Beyond establishing standards by which models of steady-state behavioral dynamics can be judged, this report also has raised a number of open questions. From an empirical viewpoint, is the breadth of the distributions seen in Figure 3 an inherent feature of repeated trials dynamics in steady state, or can the variance be decreased to allow higher precision model tests? In a related question that mixes empirical and theoretical concerns, do schedules or procedures exist that are more completely described by a linear

analysis? One can also open the range of generalization during the test phase by varying the VI during the pulse or changing the trial length. Taking a different tack, one could explore the factors that control the ringing, as shown in Figure 3, by varying the number, durations, and separations of the pulses. From a theoretical perspective, the prediction algorithm in Figure 1 is still the simplest method to extract and use a linear transfer function. How much improvement can be made in the fidelity of the predictions to measured behavior with the iterative or constrained transfer function estimation techniques? Can the generalization range be expanded to predict steady-state FI schedule behavior based on the behavioral dynamics measured with multiple VI pulses? Finally, can one broaden the analysis of steady-state repeated trials behavioral dynamics with the explicit inclusion of nonlinear terms? The questions would then become: What types of adaptive second-order correlative filters will have linear-like steady-state dynamics in repeated trials, how can one experimentally determine the properties of such filters, and what is the generalization test for filters of this class?

REFERENCES

Ayers, G. R., & Dainty, J. C. (1988). Iterative blind deconvolution method and its applications. *Optics Letters, 13,* 547–549.

Bevington, P. R. (1969). *Data reduction and error analysis for the physical sciences.* New York: McGraw-Hill.

Bracewell, R. N. (1986). *The Fourier transform and its applications* (2nd ed., rev.). New York: McGraw-Hill.

Fienup, J. R. (1993). Phase-retrieval algorithms for a complicated optical system. *Applied Optics, 32,* 1737–1746.

Fleshler, M., & Hoffman, H. S. (1962). A progression for generating variable-interval schedules. *Journal of the Experimental Analysis of Behavior, 5,* 529–530.

Higa, J. J. (1996). Dynamics of time discrimination: II. The effects of multiple impulses. *Journal of the Experimental Analysis of Behavior, 66,* 117–143.

Higa, J. J., & Pierson, D. (1998). Temporal control in rats: Analysis of nonlocalized effects from short interfood intervals. *Journal of the Experimental Analysis of Behavior, 70,* 35–43.

Horner, J. M., Staddon, J. E. R., & Lozano, K. K. (1997). Integration of reinforcement effects over time. *Animal Learning & Behavior, 25,* 84–98.

Kessel, R. (1998). Fixing a hole where the zeros get in that stop the transfer function from functioning. Retrieved from http://www.jsu.edu/depart/psychology/sebac/xfer/vi_extinc4.html.

McDowell, J. J, Bass, R., & Kessel, R. (1983). Variable-interval rate equations and reinforcement and response distributions. *Psychological Review, 90,* 364–375.

McDowell, J. J, Bass, R., & Kessel, R. (1992). Applying linear systems analysis to dynamic behavior. *Journal of the Experimental Analysis of Behavior, 57,* 377–391.

McDowell, J. J, Bass, R., & Kessel, R. (1993). A new understanding of the foundation of linear system theory and an extension to nonlinear cases. *Psychological Review, 100,* 407–419.

Palya, W. L. (1991). Laser printers as powerful tools for the scientific visualization of behavior. *Behavior Research Methods, Instruments, & Computers, 23,* 277–282.

Palya, W. L. (1992). Dynamics in the fine structure of schedule-controlled behavior. *Journal of the Experimental Analysis of Behavior, 57,* 267–278.

Palya, W. L. (1993). Bipolar control in fixed interfood intervals. *Journal of the Experimental Analysis of Behavior, 60,* 345–359.

Palya, W. L., & Walter, D. E. (1993). A powerful, inexpensive experiment controller or IBM PC interface and experiment control language. *Behavior Research Methods, Instruments, & Computers, 25,* 127–136.

Palya, W. L., Walter, D. E., Kessel, R., & Lucke, R. (1996). Investigating dynamic behavior with a fixed-time extinction schedule and linear analysis. *Journal of the Experimental Analysis of Behavior 66,* 391–409.

Press, W. H., Flannery, B. P., Teukolsky, S. A., & Vetterling, W. T. (1996). *Numerical recipes in Fortran 77: The art of scientific computing* (2nd ed.). Cambridge, England: Cambridge University Press.

Ramirez, R. W. (1985). *The FFT, fundamentals and concepts.* Englewood Cliffs, NJ: Prentice Hall.

Received June 21, 2000
Final acceptance September 29, 2001

APPENDIX

Both the data reduction and the linear dynamics prediction reuse computational methods introduced by Palya et al. (1996). Because this earlier discussion is available, we will review these methods and restate the underlying analytic definitions, but will forgo a complete step-by-step development.[4]

Given that dynamic behavior is currently understood at only a rudimentary level, we will focus on the central tendency of the data and to a lesser extent on its variance. One method to extract these properties from steady-state dynamics is to use the repeated trials local average. There are numerous ways that behavior could change as a function of time, and the use of averaging, by definition, sweeps some of these possibilities from view. Within behavior analysis the most familiar example of the repeated trials local average is

[4] A complete tutorial using an RC low-pass filter test case is available at http://www.jsu.edu/depart/psychology/sebac/low-pass/.

probably the computation of the FI scallop. The FI between reinforcer availabilities is subdivided into a set of equal-duration time bins. The number of responses in each bin is then averaged across successive FI periods to yield the scallop. In the present case, the 2,000-s trials are split into 256 bins. The number of responses in each bin is then summed across the corresponding bins from all trials across 20 sessions in a given phase. Finally, the summed number of responses in each bin is converted to a repeated trials local average response rate by dividing by both the number of trials and the constant bin period. An identical computation run on the reinforcer delivery times resulted in a repeated trials local average reinforcement rate.

The algorithm outlined in Figure 1 is useful because the dynamic relation between supported behavior and the reinforcement schedule is substantially simplified by transforming into the frequency domain. Our analysis uses Fourier transforms for the required conversions between the time and frequency domains. The forward Fourier transform,

$$h(f) = \int_{-\infty}^{\infty} H(t)\, e^{2\pi i f t}\, dt, \qquad (2)$$

converts the function of time $H(t)$ to the corresponding function of frequency $h(f)$. Conversely, a function of time can be recovered from a function of frequency with the back or inverse Fourier transform:

$$H(t) = \int_{-\infty}^{\infty} h(f)\, e^{-2\pi i f t}\, df. \qquad (3)$$

Although Equations 2 and 3 define Fourier transforms for continuous functions, the data were discretely sampled in 256 equally spaced bins. Conveniently, there is a discrete form of the Fourier transform called the Fast Fourier Transform (FFT) that is particularly well suited to computer evaluation. Our data-reduction software relies on the FFT implementation of Press, Flannery, Teukolsky, and Vetterling (1996), which is based on the definitions in Equations 2 and 3. As already noted, because an FFT yields both amplitude and phase information in the frequency domain, the computations for both transfer functions and predictions use complex arithmetic.

The use of the discretely sampled FFT and the form of the data combine to determine the frequency scale. The lowest nonzero frequency component is

$$f_1 = \frac{1}{T_{\text{trial}}}, \qquad (4)$$

where T_{trial} = 2,000 s for the present experiment. The FFT frequency scale is limited by the highest frequency component at the Nyquist critical frequency, or cut-off. The Nyquist frequency is given by

$$f_c = \frac{1}{2\Delta}, \qquad (5)$$

where Δ is the time between successive samples. In the present experiment Δ = 2,000/256 = 7.81250 s. One of the principal results from Palya et al. (1996) is that the transfer functions of pigeons are some type of low-pass filter that begins blocking frequencies above roughly 30 mHz (1 mHz is a frequency with a 1,000-s period). This characteristic 30-mHz frequency appears as a bend in transfer function magnitude when plotted in logarithmic coordinates and is normally termed the filter's "knee frequency." Consequently, we selected the frequency scale for the present experiments to bracket the knee frequency and improve the resolution below it. The 2,000-s trial length sets f_1 at 0.5 mHz, doubling the number of low-frequency components over 1,000-s trials used by Palya et al. The use of 256 samples over the 2,000-s trial results in a Nyquist cut-off at 64.0 mHz.

Once the transfer function has been determined and a prediction made for a new schedule of reinforcement, we have predicted and measured response rates for each of the 256 time bins. The quantitative question is: How close to the measured response rates are the predictions? "Close to the mark" has a somewhat more involved meaning when comparing a pair of curves than when scoring archery or darts. The reduced chi-squared sum of squared discrepancies is the standard method to answer this question (Bevington, 1969). Following Bevington, as well as Palya et al. (1996), χ^2_ν for the discrepancy between measured response rates $B(t)$ and the predicted response rates $B^{\text{pred}}(t)$ is

$$\chi_\nu^2 = \frac{1}{N - n} \sum_{i=1}^{N} \frac{1}{\sigma_{B_i}^2} [B(t_i) - B^{\text{pred}}(t_i)]^2, \quad (6)$$

$$= \frac{1}{N} \sum_{i=1}^{N} \frac{1}{\sigma_{B_i}^2} [B(t_i) - B^{\text{pred}}(t_i)]^2, \quad (7)$$

where ν, the degrees of freedom, is the difference between the number of data points N and the number of adjustable parameters n, and $\sigma_{B_i}^2$ is the variance associated with the ith response-rate residual $B(t_i) - B^{\text{pred}}(t_i)$. Because the predictions have no free parameters, $n = 0$, and the normalization factor simplifies, as shown, to $1/N$ and $\nu = N$. Note that in the more general case where $\nu < N$ and Equation 6 applies, a prediction's χ_ν^2 is penalized for free parameters through an increase in the normalization factor. In most circumstances, a prediction will not have an inherent width, and $\sigma_{B_i}^2$ reduces to just the variance in the ith measured response rate $B(t_i)$. When $B^{\text{pred}}(t_i)$ is spread over some range with a variance $\sigma_{B^{\text{pred}}}^2$, then $\sigma_{B_i}^2$ is the sum of the two contributions

$$\sigma_{B_i}^2 = \sigma_{B(t_i)}^2 + \sigma_{B^{\text{pred}}(t_i)}^2. \quad (8)$$

(To express Equation 8 in terms of standard deviations, the two contributions are added in quadrature.) Variance estimates appropriate for the current experiment are discussed in the Results section.

In our present application, Equation 7 measures how close the prediction comes to the correct order and value of the 256 measured rates. In general, the use of χ_ν^2 takes two forms: relative and absolute. Both forms will have a role in the analysis of this experiment's data. A relative χ_ν^2 comparison identifies the better of two (or more) predictions by the lower χ_ν^2 value. Provided that the same method is used to compute the $\sigma_{B_i}^2$ values, the relation between χ_ν^2 values is insensitive to the absolute magnitude of the $\sigma_{B_i}^2$ values. In contrast, absolute χ_ν^2 assesses a prediction's quality without reference to other models. Instead, one compares the residuals against a chance or random prediction scaled by the uncertainty weighting given by the $\sigma_{B_i}^2$ values. The probability q that agreement between the data and prediction did not occur by chance is given by

$$q = Q\left(\frac{\nu}{2}, \frac{\nu\chi_\nu^2}{2}\right), \quad (9)$$

where Q is the incomplete gamma function

$$Q(a, x) = \frac{\int_x^\infty e^{-t} t^{a-1}\, dt}{\int_0^\infty e^{-t} t^{a-1}\, dt} \qquad (a > 0) \quad (10)$$

(Bevington, 1969; Press et al., 1996). Strictly, a q close to 1.0 does not give the probability that a prediction is correct, but instead, the conditional probability for the agreement between observation and prediction once one assumes that the model is correct. (Of course, should some new model consistently provide better agreement of observation using relative χ_ν^2, the typical course is to abandon the presumption of correctness for the older model.) The use of Equation 9 can place unrealistic demands on the form and the reliability of the $\sigma_{B_i}^2$ values. The distribution of the residuals is implicitly assumed to be well approximated by a normal distribution. Further, uncertainty estimates for the residual between an experimentally measured quantity and its predicted value are often a matter of art. Consequently, although Equation 9 does formally answer for the goodness of fit of a prediction, rigorous mathematical certainty may exist only in situations in which the phenomenon under study is thoroughly understood.

As an aside, note that the chi-squared measure of discrepancy is also the basis for the familiar least squares line fit. The contrast between the general case, such as obtains in this report, and the linear special case is important. Because of the high symmetry of a straight line, simple explicit expressions exist for the best slope and intercept in a least squares sense. Further, the high symmetry of the line also gives rise to the correlation coefficient, r^2, which has the special property of comparing a prediction to an ideal result without explicit reference to the variance in the data. Depending on the analytic characteristics of the curve involved, similar closed-form least squares calculations are sometimes possible, although they quickly become complex and specialized. Consequently, in a more general case one normally tests the significance of a prediction using χ_ν^2 and Equation 9 directly.